



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Multi-period Trading Prediction Markets with Connections to Machine Learning

Citation for published version:

Hu, J & Storkey, A 2014, Multi-period Trading Prediction Markets with Connections to Machine Learning. in *Proceedings of ICML 2014*. vol. 32, Journal of Machine Learning Research: Workshop and Conference Proceedings, pp. 1773-1781. <<http://jmlr.org/proceedings/papers/v32/hu14.html>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Peer reviewed version

Published In:

Proceedings of ICML 2014

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Multi-period Trading Prediction Markets with Connections to Machine Learning

Jinli Hu
Amos Storkey

J.HU@ED.AC.UK
A.STORKEY@ED.AC.UK

School of Informatics, The University of Edinburgh, 10 Crichton Street, Edinburgh, EH8 9AB

Abstract

We present a new model for prediction markets, in which we use risk measures to model agents and introduce a market maker to describe the trading process. This specific choice of modelling approach enables us to show that the whole market approaches a global objective, despite the fact that the market is designed such that each agent only cares about its own goal. In addition, the market dynamic provides a sensible algorithm for optimising the global objective. An intimate connection between machine learning and our markets is thus established, such that we could 1) analyse a market by applying machine learning methods to the global objective; and 2) solve machine learning problems by setting up and running certain markets.

1. Introduction

Following the mainstream interest in “big data”, one valuable direction of machine learning is towards building up distributed, scalable and self-incentivised systems which could organise for solving large scale problems. Recently, prediction markets (Wolfers and Zitzewitz, 2004) show promise of being an abstract framework for machine learners to design these systems. As one type of markets, prediction markets naturally introduce the concepts such as self-incentivised computation and distributed environment. Additionally, the close relationship between prediction markets and probabilities sheds light on a new way of achieving probabilistic modelling (Storkey, 2011).

Since Pennock and Wellman (1996), researchers have spent decades on building connections between machine learning and prediction markets. However there is still much scope

for research in this area. One reason is that the framework of prediction markets still leaves open many design decisions, and in order to analyse the markets for machine learning goals one has to first specify a market model to describe the prediction markets. The other reason is that given a market model, we may still not know what the market is doing, even if we understand agent behaviours and market mechanisms. As distinct from most machine learning methods which explicitly define and optimise certain objectives, markets only introduce local objectives for each individual agent. To interpret a market as a machine learning method, we have to find the global objective that the market aims to optimise. This idea motivates our work.

Instead of just focusing on market mechanisms (Chen and Wortman Vaughan, 2010), we would like to incorporate the agents and analyse our market as a whole. This setting is similar to Storkey (2011); Frongillo et al. (2012); Barbu and Lay (2012); but unlike Barbu and Lay (2012), we will build a model on agent behaviours; and unlike Storkey (2011) and Frongillo et al. (2012), we model agents using risk measures, which provides analytical advantages.

The novel results of this paper include:

- a simple model for whole markets, which includes models of both the market mechanism and agents, and is easy to analyse.
- the analysis of the model which shows that there is a global objective that the market aims to optimise as a whole, and that the market trading process forms a sequential optimisation of it;
- a primal-dual relation that exists between the market and a class of machine learning problems, such that we could leverage one to solve the other.

2. A General Prediction Market Setup

Let Ω be the space of all possible future states. We say a prediction market is built on Ω if it trades securities associated with the future state $\omega \in \Omega$. Specifically, secu-

rities are defined as a set of random variables $\{\xi_k(\cdot)\} = \{\xi_1(\cdot), \xi_2(\cdot) \dots, \xi_K(\cdot)\}$. Each $\xi_k(\cdot) : \Omega \rightarrow \mathbb{R}$ is a payment function, that is, one unit of this security will pay to the holder $\xi_k(\omega)$ if ω turns out to be the future state. This definition is quite general, and securities defined in this way are also referred to as *complex securities* (Abernethy et al., 2013). We require that all securities $\{\xi_k(\cdot)\}$ (collected into the vector $\xi(\cdot)$) are linearly independent, that is, for $\mathbf{a} \in \mathbb{R}^K$, we have $\mathbf{a} \cdot \xi(\cdot) = 0$ only if $\mathbf{a} = \mathbf{0}$. If they are not, then we can always pick a subset $\{\xi_{k'}(\cdot)\}$ of linearly independent securities from $\{\xi_k(\cdot)\}$ such that all the other securities in $\{\xi_k(\cdot)\}$ can be represented by the linear combination of $\{\xi_{k'}(\cdot)\}$ (Kreyszig, 2007). Therefore it is redundant to consider $\{\xi_k(\cdot)\}$ that are not linearly independent. As an example, the Arrow-Debreu security is a special case of complex securities. When the sample space Ω is discrete and contains only finite number of states, Arrow-Debreu securities are a set of $K = |\Omega|$ securities, in which the k -th one pays one unit if the k -th state is true: $\xi_k(\omega) = \mathbb{1}(\omega = k)$. Note that in general cases $K < |\Omega|$, e.g. when the value of ω is continuous, there will be infinite number of states but we always have a finite K for practice.

Agents can only trade these predefined securities. The behaviour of an agent is characterised by its *portfolio* $\{w, s_k\} = \{w, s_1, s_2, \dots, s_K\}$, where w is the amount of money that the agent has, and s_k is the amount of shares the agent holds in security k . We collect all s_k into vector \mathbf{s} . If an agent has a portfolio $\{w, s_k\}$, the total payment of the securities is

$$X(\cdot) = \mathbf{s} \cdot \xi(\cdot), \quad (1)$$

where $X(\cdot) : \Omega \rightarrow \mathbb{R}$ is in essence a random variable on Ω . We call $X(\cdot)$ the *risky asset* because of its uncertainty and w the *risk-free asset*. The gross payment is thus

$$\hat{X}(\cdot) = w + \mathbf{s} \cdot \xi(\cdot) = w + X(\cdot), \quad (2)$$

which is also a random variable. We call $\hat{X}(\cdot)$ the (*gross*) *asset*. Denote \mathcal{X} the set of all $X(\cdot)$ that are accessible for an agent, and similarly $\hat{\mathcal{X}}$ the set of all $\hat{X}(\cdot)$. Notice that since $\{\xi_k(\cdot)\}$ are linearly independent, there exists a unique map (bijection) between $X(\cdot)$ and \mathbf{s} via (1). Therefore a portfolio could also be represented by $\{w, X(\cdot)\}$. In our setting $\mathcal{X} \subseteq \text{span}(\xi_1(\cdot), \xi_2(\cdot) \dots, \xi_K(\cdot))$, but it is possible to make \mathcal{X} more abstract, which is not a space spanned by a prefixed number of securities but allows new security types to be added on the fly. This discussion is beyond our scope.

A market consists of two processes, that 1) each agent chooses a portfolio $\{w, X(\cdot)\}$ it would like to hold, and 2) agents try to move to their preferred portfolio by trading. To describe the decision making process we need a model of portfolio selection, while to describe trading we need to specify a market mechanism. These two parts will be discussed in Section 3 and 4.

Later in this paper, when the context is clear we will omit parentheses and write a random variable in an uppercase letter, e.g. X (except the securities, which are denoted by ξ), and use the lowercase of the same letter for the value of it, e.g. x . We will also write functionals in letters without any parentheses.

3. Preferences on Assets

Agents select assets based on their preferences. An agent's preference order of two assets is measured by a functional $f : \hat{\mathcal{X}} \rightarrow \mathbb{R}$, such that the agent prefers one asset \hat{X} than the other asset \hat{Y} if and only if $f(\hat{X}) > f(\hat{Y})$, and that the agent is indifferent between X and Y if and only if $f(\hat{X}) = f(\hat{Y})$. There are plenty of theories on choosing and analysing a specific form of f . These includes expected utility theory (EUT) (Von Neumann and Morgenstern, 2007), dual utility theory (Yaari, 1987), risk measures (Artzner et al., 1999), etc. EUT is perhaps the most popular theories in economics and game theory, while risk measures are commonly seen in finance literature. We choose to use risk measures to model agent behaviours. We introduce risk measures in this section, while putting the detailed justification of using risk measures and its relation to EUT in Section 6.

3.1. Risk measures

As is indicated by their name, risk measures assign higher scores to assets that are more "risky". They can also be understood as measures of the potential loss of choosing certain asset. A (*monetary*) *risk measure* is defined as a functional $\rho : \mathcal{X} \rightarrow \mathbb{R}$ such that $\rho(0)$ is finite and ρ satisfies the following conditions (Artzner et al., 1999):

Translation invariance If $X \in \mathcal{X}$ and $m \in \mathbb{R}$, then

$$\rho(X + m) = \rho(X) - m. \quad (3)$$

Monotonicity If $X, Y \in \mathcal{X}$ and $X \leq Y$, then

$$\rho(X) \geq \rho(Y). \quad (4)$$

Here $X \leq Y$ should be understood as $P(x \leq y) = 1$, that is, with the probability of one that X will generate a lower return than Y . Thus monotonicity indicates that an asset with a better return deserves a lower risk. Due to translation invariance, a risk measure maps any risk-free asset to itself, and is additive w.r.t. any risk-free asset. Therefore, the output of a risk measure has the same unit with a risk-free asset, and can be calculated like an asset.

The domains of risk measures and the preference functional f are different, as risk measures are defined on \mathcal{X} while the space of assets that agent can hold is $\hat{\mathcal{X}}$. Fortunately, we

could easily extend the definition of risk measures to the domain $\hat{\mathcal{X}}$ by applying translation invariance (3)

$$\rho(\hat{X}) = \rho(X + w) = \rho(X) - w, \quad \forall \hat{X} \in \hat{\mathcal{X}}. \quad (5)$$

A corresponding f can thus be obtained by $f = -\rho$.

Risk measures are very generic. In our discussion we will use both risk measures and a specific class of them, the *convex* risk measures (Föllmer and Schied, 2002). A risk measure is convex if $\forall X_1, X_2 \in \mathcal{X}$ and $\lambda \in [0, 1]$

$$\rho(\lambda X_1 + (1 - \lambda)X_2) \leq \lambda \rho(X_1) + (1 - \lambda)\rho(X_2). \quad (6)$$

It says that the risk of a combination of two assets should not be higher than holding them separately. In other words, convex risk measures encourage diversification, which is a natural condition on risk measures.

Examples of risk measures A famous non-convex risk measure is the *Value at Risk* (VaR) (Linsmeier and Pearson, 2000), which outputs a threshold loss l such that the probability of $-X$ exceeding l is smaller than a predefined level

$$\text{VaR}_\alpha(X) \equiv \inf\{l \in \mathbb{R} \mid P(-X > l) \leq 1 - \alpha\}. \quad (7)$$

A famous convex risk measure is the *Entropic risk measure* (Föllmer and Schied, 2004)

$$\rho_E = \frac{1}{\theta} \log M_X(-\theta) = \sup_{Q \in \mathcal{P}} \mathbb{E}_Q[-X] - \frac{1}{\theta} D[Q \| P]. \quad (8)$$

Here $M_X(t) \equiv \mathbb{E}_P[e^{tX}]$ is the moment-generating function, and $D[\cdot \| \cdot]$ is the KL-divergence (and this is where “entropic” comes in). We mention that the second representation of ρ_E holds for all convex risk measures, and this representation becomes the key to connecting the markets to machine learning (cf. Section 5).

3.2. Rational Choices

Recall that a portfolio that leads to a higher value of $f(\hat{X})$ is preferred. Thus the favourite portfolio of an agent should be the one that maximises f , which we denote by $\{w, X\}^{opt}$. The behaviour of choosing $\{w, X\}^{opt}$ is called the *rational choice*, and an agent is rational if it always chooses $\{w, X\}^{opt}$ as its trading goal. Since in our framework $f = -\rho$, a rational agent will choose will $\{w, X\}^{opt}$ under the rule of

$$\min_{\{w, X\}} \rho(\hat{X}) = \min_{\{w, X\}} \rho(w + X). \quad (9)$$

In a market an agent only cares about its own goal (9). It seems like this property prevents us from linking markets to machine learning methods, as the latter always aim to achieve certain global objectives. However, with a careful design, we can let our markets implicitly define global objectives and make an agent contribute to the global objective at the same time as it achieves its own goal.

4. Multi-period Trading Markets

In this section we will build our market, a multi-period trading market whose trades are driven by a market maker. “Multi-period” is used to indicate that the prices of the securities are allowed to vary at different time steps, and that agents can trade with the market maker at multiple times (Föllmer and Schied, 2004). The market maker is introduced to simplify the market mechanism and to make the market run efficiently.

It is difficult to characterise the trading process in the markets with unspecified mechanisms, and those markets may not run efficiently. For example, there may not exist a consistent agreement among agents on how much should be paid to buy/sell one share of a security. Moreover, one agent who wants to sell a certain amount of shares may not find any buyers (Chen and Pennock, 2007). One way to simplify the trading process is by introducing a market maker (Hanson, 2007). A market maker is a special agent. It is a price maker, who defines the price for trading each security. All agents are only allowed to trade with the market maker. They can, however, make a trade at any time as long as they agree to pay under the market maker’s pricing. The pricing rule of a market maker at time step t is a functional $c_t : \mathcal{X} \rightarrow \mathbb{R}$. At different time steps the cost for purchasing an asset may be different, i.e. it may happen that $c_t(X) \neq c_{t'}(X)$ when $t \neq t'$.

Suppose that an agent has a portfolio $\{w_{t-1}, X_{t-1}\}$ at time $t - 1$ and it would like to buy ΔX_t from the market maker at t . The agent cannot propose an arbitrary price Δw_t for ΔX_t but has to accept the price provided by the market maker $\Delta w_t = -c_t(\Delta X_t)$. The updated portfolio is thus restricted to $\{w_{t-1} - c_t(\Delta X_t), X_{t-1} + \Delta X_t\}$, and the updated asset is restricted to $\hat{X}_t = X_{t-1} + w_{t-1} + \Delta X_t - c_t(\Delta X_t)$. Now a rational agent only cares about choosing its optimal purchase amount ΔX_t such that $\rho(\hat{X}_t)$ is minimised:

$$\begin{aligned} & \min_{\{w_t, X_t\}} \rho(\hat{X}_t) \\ & = \min_{\Delta X_t \in \mathcal{X}} \rho(X_{t-1} + \Delta X_t + w_{t-1} - c_t(\Delta X_t)). \end{aligned} \quad (10)$$

This portfolio selection procedure leads to Algorithm 1.

Algorithm 1 **Select**($\{w_{t-1}, X_{t-1}\}, \rho(\cdot), c_t(\cdot)$): portfolio selection of a rational agent

Input: portfolio $\{w_{t-1}, X_{t-1}\}$, risk measure $\rho(\cdot)$, pricing rule $c_t(\cdot)$

Choose the ΔX_t that minimise (10)

Output: $\{\Delta X_t, -c_t(\Delta X_t)\}$

We now consider a multi-period market which involves a set $A = \{1, 2, \dots, N\}$ of agents and a market maker. Assume that at each round t there is only one agent $a_t \in A$

that trades with the market maker. This assumption indicates that each agent trades with the market maker separately, and they do not cooperate to make a joint purchase. $\{a_1, a_2, \dots, a_T\}$ is thus the trading queue of the market. Since there are multiple agents, we use an extra subscript to distinguish the portfolios of different agents. For example, an agent $n \in A$'s portfolio at time t is $\{w_{n,t}, X_{n,t}\}$. The initial values are denoted with the subscript $t = 0$. We collect all $w_{n,t}, X_{n,t}, \hat{X}_{n,t}$ into vectors $\mathbf{w}_t, \mathbf{X}_t$ and $\hat{\mathbf{X}}_t$, respectively. We assume that agents do not bring in any risky asset at the beginning, which is a natural assumption since only the market maker can issue securities. This assumption means we have $\mathbf{X}_0 = \mathbf{0}$ and so $\hat{\mathbf{X}}_0 = \mathbf{w}_0$.

At time t , only the agent a_t updates its portfolio by trading with the market maker while all the other agents keep the same portfolios as at $t - 1$. Suppose the asset that the agent a_t would like to purchase is $\Delta X_{a_t,t}$, then for all $n \in A$

$$X_{n,t} = X_{n,t-1} + \mathbb{1}(n = a_t) \Delta X_{a_t,t}, \quad (11a)$$

$$w_{n,t} = w_{n,t-1} - \mathbb{1}(n = a_t) c_t(\Delta X_{a_t,t}). \quad (11b)$$

Algorithm 2 runs a multi-period trading market.

Algorithm 2 A multi-period market with a set A of agents and a market maker

Input: initial portfolios $\{\mathbf{w}, \mathbf{X}_0\}$, risk measures $\{\rho_n(\cdot)\}$, pricing rule $c_t(\cdot)$, time period T

for $t = 1$ **to** T **do**

for each agent $n \in A$ **do**

 propose $\{\Delta X_{n,t}, -c_t(\Delta X_t)\}$ using Algorithm 1

end for

 trade happens between the market maker and a_t

for each agent $n \in A$ **do**

 update their portfolios using (11)

end for

end for

We can also split Algorithm 2 into the market maker routine and the agent routine (details in supp.). We do this to emphasise the fact that each agent has its *own objective* (achieving the optimal portfolio based on its unique preference), plus a communication with the market maker.

4.1. Appropriate choice of the pricing rule $c_t(\cdot)$

There has been plenty of work on studying the pricing rule $c_t(\cdot)$ of a market maker (Brahma et al., 2012; Pennock, 2004). A popular class of mechanisms is Hanson's market scoring rules (Hanson, 2007). It is later formalised by Abernethy et al. (2013), who use a set of reasonable axioms to characterise the pricing mechanism. We apply their result to our framework.

Let $\Delta X_t \equiv \Delta X_{a_t,t}$ be the trade with the market maker at time t . Consider two situations: 1) a trade happens

with the market maker in ΔX ; and 2) a trade happens with the market maker in $\Delta X'$ and is followed by another trade $\Delta X''$, where $\Delta X = \Delta X' + \Delta X''$. A natural requirement is that the cost of purchasing ΔX should be equal to the total cost of purchasing $\Delta X'$ and $\Delta X''$. Under this condition, Abernethy et al. (2013) show that there exists a functional $c : \mathcal{X} \rightarrow \mathbb{R}$ which has the form

$$c_t(\Delta X_t) = c(\Delta X_1 + \dots + \Delta X_{t-1} + \Delta X_t) - c(\Delta X_1 + \dots + \Delta X_{t-1}). \quad (12)$$

We say a pricing rule c_t is *path-independent* if it has the form of (12), and reload the notation c to represent c_t .

5. The Machine Learning Objective of the Multi-period Trading Markets

The primary goal of this paper is to establish an intimate connection between machine learning and our new prediction market model. Before we start to analyse the multi-period trading markets, we introduce the machine learning context for which we want our markets to be utilised. Many machine learning tasks could be interpreted under the following generic framework: given a set of data sampled from a space Ω , and a hypothesis space \mathcal{P} which contains a class of accessible probabilities on Ω , we would like to find a probability from \mathcal{P} that can best describe the data. Usually we use a functional $F : \mathcal{P} \rightarrow \mathbb{R}$ to characterise the "best" performance, such that the best probability is the one that minimises F . Formally, this involves an optimisation problem

$$\min_{P \in \mathcal{P}} F(P) \quad (13)$$

For specific problems in which the information comes from different parts of the data or the models, F has a form of $F = \sum_n F_n$, the sum of a set of functionals which share the same domain \mathcal{P} (see examples in Section 7 for details). We will show that a multi-period market effectively defines and optimises a machine learning task whose $F = \sum_n F_n(P)$.

The connection is established in two steps: first we show that the market does have a global objective, and then show that under mild conditions the market optimises the dual of a machine learning problem $\min_{P \in \mathcal{P}} \sum_n F_n$.

5.1. The global objective of a market

We show that a multi-period trading market minimises a global objective. The optimisation is done sequentially via the market trading dynamics, that is, an agent will contribute to minimising this global objective as long as it makes a trade with the market maker. This argument is formalised in the following

Proposition 1 (The global objective of a market). *A multi-period market (Algorithm 2) with a path-independent pricing rule market maker aims to minimise the global objective*

$$L = c(Y) + \sum_{n \in A} \rho_n(X_n), \quad Y = \sum_{n \in A} X_n, \quad (14)$$

by performing a sequential optimisation algorithm, which is implemented by the market trading process (cf. (10) and (11)): define $\varphi_{a_t}(\Delta X'_t) \equiv \rho_{a_t}(X_{a_t,t-1} + \Delta X'_t + w_{a_t,t-1} - c_t(\Delta X'_t))$ and for each time t

$$\Delta X_t = \arg \min_{\Delta X'_t} \varphi_{a_t}(\Delta X'_t), \quad (15a)$$

$$X_{n,t} = X_{n,t-1} + \mathbb{1}(n = a_t) \Delta X_t, \quad (15b)$$

$$w_{n,t} = w_{n,t-1} - \mathbb{1}(n = a_t) c_t(\Delta X_t), \quad (15c)$$

$$Y_t = Y_{t-1} + \Delta X_t, \quad (15d)$$

If the algorithm converges at time t' , i.e. $\Delta X_t = 0$ for all $t > t'$, then $\{X_{n,t'}\}, Y_{t'}$ achieves a local minimum of the objective L in (14).

Proof. Outline (details in supp.): recall that at time t only agent a_t will trade with the market maker, so $\Delta X_t = \Delta X_{a_t,t}$ and $\Delta X_{n,t} = 0, \forall n \neq a_t$. At time t , for any agent n all quantities calculated before t can be treated as constants as they could no longer be modified. Therefore, the functional that is minimised in (15a) has the same optimal point as the following functional

$$l_t(\Delta X'_t) = \rho_{a_t}(X_{a_t,t-1} + \Delta X'_t + w_{a_t,t-1} - c_t(\Delta X'_t)) - \rho_{a_t}(X_{a_t,t-1} + w_{a_t,t-1}). \quad (16)$$

Define $L_T = \sum_{t=1}^T l_t$ and use the translation invariance of a risk measure and the path-independence of the pricing rule. We will end up with

$$\min_{\{\Delta X_t\}} L_T = \min_{\{Y_T\}} c(Y_T) + \sum_{n \in A} \rho_n(X_{n,T}) - C, \quad (17)$$

where $Y_t = \sum_{\tau=1}^t \Delta X_\tau = \sum_{n \in A} X_{n,t}$ holds for $\forall t > 0$, and C is a constant. (17) is a sequential minimisation scheme for $\min L$. Finally, if the market converges at time T , we have $X_n = X_{n,T}$ and $Y = Y_T$, leading to a local minimal point of L . \square

Proposition 1 is the key to understanding the market mechanism. Despite the fact that the market is set up to let agents behave under their own preferences, the market mechanism ensures that a global objective is established, and that the agent will contribute to optimising the global objective at the same time as it optimise its own goal. The trading process thus provides a sensible algorithm for achieving this global objective.

5.2. A primal-dual representation via convex analysis

One concern is that (14) is not commonly seen in machine learning problems¹. A different view of this objective should somehow be introduced. In fact, under mild requirements on the form of risk measures and pricing rules, the global objective forms the dual of the optimisation problem $\min_{P \in \mathcal{P}} \sum_n F_n(P)$. The requirement for the risk measures is convexity (6). The requirement for the pricing rules is that it is duality-based (Abernethy et al., 2013).

5.2.1. MORE ON CONVEX RISK MEASURES

Artzner et al. (1999) and Föllmer and Schied (2002) show that a convex risk measure has a form

$$\rho(X) = \sup_{Q \in \mathcal{P}} (\mathbb{E}_Q[-X] - \alpha(Q)), \quad (18)$$

where \mathcal{P} is a set of probabilities on (Ω, \mathfrak{F}) such that Q is absolutely continuous w.r.t. P and $\mathbb{E}_Q[X]$ is well defined. The risk measure decreases as $\mathbb{E}_Q[X]$ increases but this effect is penalised by a functional α . (18) is in essence a Legendre-Fenchel transform with a slight change on signs (Boyd and Vandenberghe, 2004).

5.2.2. DUALITY-BASED PRICING RULES

We keep following the idea of Abernethy et al. (2013) and apply their duality-based pricing rules to our problem. The authors point out that duality-based pricing rules are well motivated as they meet some natural conditions such as no-arbitrage. A duality-based pricing rule is path-independent and has a form²

$$c(X) \equiv \sup_{Q \in \mathcal{P}} (\mathbb{E}_Q[X] - R(Q)) = R^*(X), \quad (19)$$

where R^* denotes the Legendre-Fenchel transform of R . Note that in their work R is required to be convex, but this condition could be relaxed since for any R we could define $R' \equiv (R^*)^* = c^*$ to replace R , as R' is always convex (as it is a conjugate dual) and $c = (R')^* = R^*$.

5.2.3. THE PRIMAL PROBLEM

Now we are ready to show

Proposition 2 (The primal problem). *For a multi-period market which involves agents who use convex risk measures in (18) and a duality-based pricing rule market maker in (19), its global objective is a weak dual of*

$$\min_{P \in \mathcal{P}} \sum_{n=0}^N F_n(P), \quad (20)$$

¹However, to complete our discussion, we show one example that uses (14) in Section 7

²Abernethy et al. (2013) represent markets in securities $\{\xi_k\}$ and shares $\{s_k\}$. To be consistent with our framework we change the representation to assets X (cf. Section 2).

where F_0 and F_n are functionals that share the same domain \mathcal{P} . Specifically, $F_0 = R$ in (19), and $F_n = \alpha_n$ where α_n is the penalty functional of agent n .

Proof. We use the generalised Fenchel's duality (Shalev-Shwartz and Singer, 2007) to derive the Lagrange dual problem of (20). Under the generalised Fenchel's duality, the dual problem (weak duality) is

$$-\min_{\{X_n\} \in \mathcal{X}} F_0^*(Y) + \sum_{n=1}^N F_n^*(-X_n), \quad Y = \sum_{n=1}^N X_n, \quad (21)$$

where F_n^* denotes the Legendre-Fenchel transform.

We construct the convex risk measure for each agent n . use (18) and choose $\alpha = F_n$

$$\rho_n(X) = \sup_{Q \in \mathcal{P}} (\mathbb{E}_Q[-X] - F_n(Q)) = F_n^*(-X). \quad (22)$$

For the pricing rule (19) we choose $R = F_0$ and obtain $c = F_0^*$. Substitute them back to the dual problem (21) and we end up with

$$-\min_{\{X_n\}} L = -\min_{\{X_n\} \in \mathcal{X}} c(Y) + \sum_{n=1}^N \rho_n(X_n). \quad (23)$$

This matches the global objective L (cf. (14)) with a different sign. The negation sign is necessary because the Lagrange dual *lower bounds* the primal in general

$$-\min_{\{X_n\}} L \leq \min_{P \in \mathcal{P}} \sum_{n=0}^N F_n(X). \quad (24)$$

If strong duality holds (Boyd and Vandenberghe, 2004), equality holds in (24) and the global objective is the equivalent to the primal machine learning problem. \square

Proposition 2 gives us two ways of building the connection between markets and machine learning: 1) If we model a market using our framework, we could then figure out the global objective of the market and then the primal problem, which can be solved using machine learning methods. 2) More interestingly, given a machine learning problem of form (20), we could transform it to a market and solve the problem by running the market, during which we could take the advantage of some market properties, such as distributed environment and privacy, to gain extra benefits.

6. Related Work

The idea of building models for prediction markets and discussing their relation to optimisation is not novel, and significant progress has been achieved in the past few years. We will discuss the work that is closely related to ours.

In Chen and Wortman Vaughan (2010), the authors show that scoring rule market makers perform online no-regret learning. Their study focuses on the market makers while agents are not directly modelled, which motivates a framework for the whole market.

Storkey (2011) defines and analyses a type of prediction markets based on definitions on the markets, securities, and agents. Agents are modelled by as maximisers of expected utility. By analysing the equilibrium status of the market the author shows that the market can aggregate beliefs from agents to output a probability distribution over the future events. The author focuses on equilibria rather than precise market mechanisms, and does not provide any global objective of the market, which makes it difficult to link these markets to optimisation procedures.

Frongillo et al. (2012) apply market scoring rules as the market mechanism to the framework of Storkey (2011). The work shows that with a large population of agents whose portfolios are drawn from a demand distribution, the whole market implements *stochastic mirror descent*. One concern is that they suggest using EUT to model agents but they do not use it to solve the optimal portfolios for the agents. This problem is partially solved by Premachandra and Reid (2013), who derives the solution for a certain type of expected utilities. A similar setting is also studied by Sethi and Wortman Vaughan (2013). They focus more on the convergence of the market dynamics, and show how markets can aggregate beliefs by using numerical evidences.

6.1. Risk measures and EUT

Here we justify the choice of risk measures as the agent decision rules. First, the output value of a risk measure can be treated as a risk-free asset and standard linear operations are well defined for it. In comparison, an expected utility outputs a number that only has abstract meaning, i.e. to measure the degree of agent's satisfaction. Additionally, risk measures force translation invariance by definition, but expected utility functions do not have this property in general. With the help of translation invariance, the wealth w can always be separated from the risky asset X , which implies that the optimal portfolio does not depend on w . This saves us from the trouble of associating w with the aggregation weights, as the relationship between them is highly inconsistent and varies dramatically under different utilities (Storkey et al., 2012). Finally, we could always derive convex a risk measure ρ_u from any expected utility (Föllmer and Schied, 2004)

$$\rho_u(X) \equiv \inf\{m \in \mathbb{R} \mid \mathbb{E}_P[u(X + m)] \geq u_0\}, \quad (25)$$

where P is the personal belief of the agent. In fact, the output of this risk measure is the *risk premium*, the least

amount of money that one would like to borrow in order to accept this risky asset. Then a sensible decision rule should be to find an asset that minimises the premium, which leads to risk minimisation. For example, the entropic risk measure in (8) can be given by the exponential utility $u = -\exp(-ax)$, with $\theta = a$ (Föllmer and Schied, 2002).

7. Examples

In this section we use three examples to illustrate the connections between the multi-period trading markets and machine learning.

Opinion Pooling The opinion pooling problem is a common setting for prediction market models (Barbu and Lay, 2012; Storkey et al., 2012). Garg et al. (2004) show that the objective of an opinion pool is to minimise a weighted sum of a set of divergences. Particularly, for logarithmic opinion pool the objective is to

$$\min_{P \in \mathcal{P}} \sum_n w_n D[P||P_n]. \quad (26)$$

where $D[\cdot||\cdot]$ is the KL-divergence and $\{w_n\}$ are weight parameters.

Now consider an log-opinion pool of a set of A probabilities on a finite discrete sample space Ω with K states. To set up a market that matches the log-opinion pool, we first define a market on the same space Ω and introduce K Arrow-Debreu securities. We introduce N agents, and assign a unique probability $P_n \in A$ to agent n as its personal belief. According to (8), agent n 's risk measure has the form

$$\rho_n(\mathbf{s}_n) = \frac{1}{\theta_n} \log \sum_{k=1}^K p_k e^{-\theta_n s_{n,k}}, \quad (27)$$

where we let θ_n match the weight w_n by $\theta_n^{-1} = w_n$. For the sake of simplicity, we choose a logarithmic market scoring rule market maker (Hanson, 2007)

$$c(\mathbf{s}_0) = \frac{1}{\theta_0} \log \sum_{k=1}^K e^{\theta_0 s_{0,k}}. \quad (28)$$

The market can be run by using Algorithm 2. Two typical simulation results are shown in Figure 1 and 2. The primal problem of this market is (applying Proposition 1 and 2)

$$\min_{P \in \mathcal{P}} \frac{1}{\theta_0} D[P||P_0] + \sum_{n \in A} \frac{1}{\theta_n} D[P||P_n], \quad (29)$$

where the domain $\mathcal{P} = \Delta_K$ is the probability simplex in K dimensions and $P_0 = \text{uniform}(K)$ is the discrete uniform distribution in Δ_K . In this case the optimal P can be analytically solved. Recall that $\theta_n^{-1} = w_n$ and we have

$$P \propto \prod_{n \in A} P_n^{w_n / (\theta_0^{-1} + \sum_{n \in A} w_n)}. \quad (30)$$

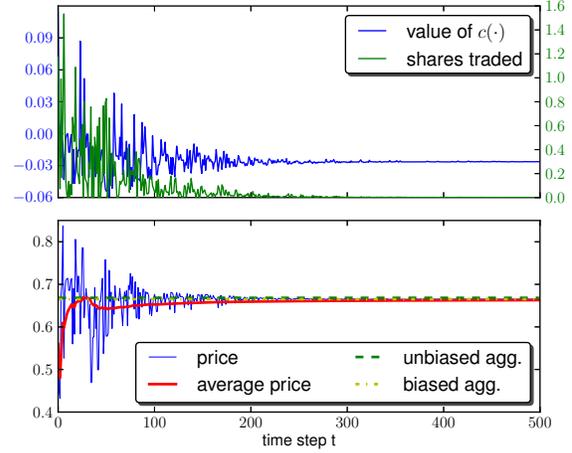


Figure 1. A market with Arrow-Debreu securities defined on a binary event ω . $N = 10$ agents are involved. All agents start with a uniform prior on ω and each one builds its own posterior belief after a private observation of 5 samples of ω . The market price converges to a position which is close to the unbiased agent aggregation, but with a small bias towards 0.5. The bias is introduced by the market maker (cf. (30)).

Since we introduce the market maker, the aggregated belief P is not a pure weighted product of agents' beliefs, but with a bias towards P_0 . However, when the population is sufficiently large such that $\sum_n \theta_n^{-1} \gg \theta_0^{-1}$, the effect of the market maker could be ignored and we will end up with a pure aggregation of agent beliefs (Frongillo et al., 2012).

Bayesian Update We give our second example by first setting up a market and then match a machine learning problem to the market. Let us build a market on a continuous sample space $\Omega = \mathbb{R}$. We only define one security $\xi(\omega) = \omega$, and so the asset $X = s\omega$. We introduce only one agent. Again, the agent is characterised by an entropic risk measures, with coefficient θ_1 and $P_1 = \mathcal{N}(\mu_1, \sigma_1^2)$ is the normal distribution. The moment-generating function in (8) is

$$M_X(-\theta_1) = \mathbb{E}_{P_1}[e^{-\theta_1 s \omega}] = e^{-\theta_1 s \mu_1 + \sigma_1^2 \theta_1^2 s^2 / 2}, \quad (31)$$

and so the risk measure is $\rho_1(s) = -s\mu_1 + \sigma_1^2 \theta_1^2 s^2 / 2$. For the market maker we use the quadratic market scoring rule $c(s) = \theta_0 s^2 / 2$. Now we could run this market using Algorithm 2 with only one agent.

It can be shown that this market implements a Bayesian *maximum a posteriori* (MAP) update for the Gaussian, in which the prior is provided by the market maker and the likelihood information is provided by the agent. The MAP update in the primal form is

$$\min_{\mu \in \mathbb{R}} \frac{1}{\theta_0} \frac{\mu^2}{2} + \frac{1}{\theta_1} \frac{(\mu - \mu_1)^2}{2\sigma_1^2}, \quad (32)$$

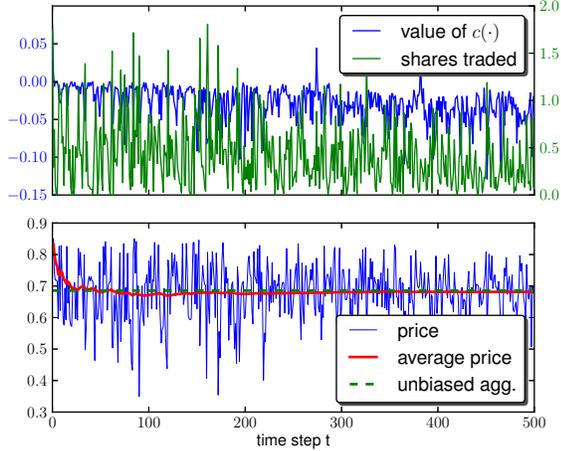


Figure 2. A market which has the same setting with Figure 1 but this time $N = 100$ agents are involved. After increasing the population, the market price does not show a sign of convergence before $t = 500$. Comparatively, the average price quickly converges to the aggregation belief. This is expected, as for a large population the market should reproduce the results of Frongillo et al. (2012).

while this update is done by the market in the dual space of the space of the mean parameter μ (details in supp.).

Logistic Regression In the third example we discuss a classical machine learning problem. Given a data set $\mathcal{D} = \{\{\mathbf{x}_m, \mathbf{y}_m\} | \mathbf{x}_m \in \mathbb{R}^K, \mathbf{y}_m = \{+1, -1\}, m = 1, \dots, M\}$, we would like to build logistic regression model with l_2 -regularisation. The objective is

$$L = \min_{\mathbf{w} \in \mathbb{R}^K} \frac{1}{M} \sum_{m=1}^M \log \left(1 + e^{y_m(\mathbf{w} \cdot \mathbf{x}_m)} \right) + \frac{\lambda}{2} \|\mathbf{w}\|^2, \quad (33)$$

where $\|\cdot\|$ is the l_2 norm.

To convert this problem to a market we use (14) and Proposition 1. Let the sample space be the space that generates the data $\Omega \equiv \mathbb{R}^K \cup \{+1, -1\}$ and each future state is associated with a data in Ω , $\omega = \{\mathbf{x}, y\}$. Define K securities, each of which is $\xi_k(\omega) = yx_k$. We introduce $N = K$ agents, such that the agent $n = k$ is only interested in trading in the k -th security ξ_k . Thus the shares of security k held by agent n is $s_{n,k} = \mathbb{1}(n = k)w_k$, and the asset is $X_n = s_n \cdot \xi = w_n \xi_n$. The market inventory is $\mathbf{s}_0 = \sum_n s_n = \mathbf{w}$. Let $c(\mathbf{w})$ be the first term on the RHS of (33) and define the risk measure of agent n as $\rho_n(\mathbf{s}_n) = \lambda s_n^2 / 2$. We end up with

$$L = \min_{\mathbf{w}} c(\mathbf{w}) + \sum_{k=1}^K \frac{\lambda}{2} w_k^2 = \min_{\{\mathbf{s}_n\}} c(\mathbf{s}_0) + \sum_{n=1}^N \rho_n(\mathbf{s}_n). \quad (34)$$

Now the market is ready to run under Algorithm 2. In order to show a slightly deeper connection to a specific learning method, we notice that the objective of agent n at each round is $\min_{\Delta w_{k,t}} c(\mathbf{w}_{t-1} + \Delta w_{k,t}) + (w_{k,t-1} + \Delta w_{k,t})^2 / 2$. As the solution to this is not analytic, it is costly to solve for the exactly minimum of this objective at each time. To get rid of this problem, we could relax the condition that agents behaviour is rationally optimal, and let the agents accept a portfolio as long as it is better than its current position $\rho_n(\hat{\mathbf{s}}_{n,t}) < \rho_n(\hat{\mathbf{s}}_{n,t-1})$. Specifically agents can take steps towards the optimal solution. This can be achieved by the following portfolio updating rule

$$\Delta w_{k,t} = -a \frac{d}{dw_k} \left(c(\mathbf{w}) + \frac{\lambda}{2} w_k^2 \right) \Big|_{\mathbf{w}=\mathbf{w}_{t-1}}, \quad (35)$$

where $a > 0$ is adjusted such that $\rho_n(\hat{\mathbf{s}}_{n,t}) < \rho_n(\hat{\mathbf{s}}_{n,t-1})$. In practice a could be chosen by backtracking line search (Boyd and Vandenberghe, 2004). The market we designed above effectively implements a *coordinate descent* algorithm (Luo and Tseng, 1992).

Note that, instead of introducing $N = K$ agents, we can match the logistic regression problem by using only one agent and allowing it to trade all securities. This will result in a standard gradient descent method.

8. Conclusion

This paper establishes and discusses a new model for prediction markets. We use risk measures instead of expected utility to model agents, which results in an analytical market framework. We show that our market as a whole optimises certain global objective through its market dynamics. Based on this result, we make intimate connections between machine learning and markets.

One area of future work would be conducting a detailed analysis of this framework using the tools of convex optimisation. A particularly interesting topic is to find the conditions under which the market will converge. As we have observed, stochasticity plays a key part when a large population of agents are involved, as is the case in most real market settings (Frongillo et al., 2012).

Acknowledgement

This work was supported by Microsoft Research Cambridge through its PhD Scholarship Programme.

References

Abernethy, J., Chen, Y., and Wortman Vaughan, J. (2013). Efficient market making via convex optimization, and a connection to online learning. *ACM Trans. Econ. Comput.*, 1(2):12:1–12:39.

- Artzner, P., Delbaen, F., Eber, J., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3):203–228.
- Barbu, A. and Lay, N. (2012). An introduction to artificial prediction markets for classification. *Journal of Machine Learning Research*, 13:2177–2204.
- Boyd, S. P. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.
- Brahma, A., Chakraborty, M., Das, S., Lavoie, A., and Magdon-Ismael, M. (2012). A Bayesian market maker. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, EC '12, pages 215–232, New York, NY, USA. ACM.
- Chen, Y. and Pennock, D. (2007). A utility framework for bounded-loss market makers. In *Proceedings of the Twenty-Third Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-07)*, pages 49–56, Corvallis, Oregon. AUAI Press.
- Chen, Y. and Wortman Vaughan, J. (2010). A new understanding of prediction markets via no-regret learning. In *Proceedings of the 11th ACM conference on Electronic commerce*, EC '10, pages 189–198, New York, NY, USA. ACM.
- Föllmer, H. and Schied, A. (2002). Convex measures of risk and trading constraints. *Finance and Stochastics*, 6(4):429–447.
- Föllmer, H. and Schied, A. (2004). *Stochastic finance: An introduction in discrete time*. Springer, 2nd revised edition edition.
- Frongillo, R. M., Penna, N. D., and Reid, M. D. (2012). Interpreting prediction markets: a stochastic approach. In Pereira, F., Burges, C., Bottou, L., and Weinberger, K., editors, *Advances in Neural Information Processing Systems 25*, pages 3275–3283.
- Garg, A., Jayram, T., Vaithyanathan, S., and Zhu, H. (2004). Generalized opinion pooling. In *AMAI*.
- Hanson, R. (2007). Logarithmic market scoring rules for modular combinatorial information aggregation. *The Journal of Prediction Markets*, 1(1):3–15.
- Kreyszig, E. (2007). *Introductory functional analysis with applications*. Wiley.
- Linsmeier, T. J. and Pearson, N. D. (2000). Value at risk. *Financial Analysts Journal*, 56(2):pp. 47–67.
- Luo, Z. and Tseng, P. (1992). On the convergence of the coordinate descent method for convex differentiable minimization. *Journal of Optimization Theory and Applications*, 72(1):7–35.
- Pennock, D. M. (2004). A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In *Proceedings of the 5th ACM conference on Electronic commerce*, EC '04, pages 170–179, New York, NY, USA. ACM.
- Pennock, D. M. and Wellman, M. P. (1996). Toward a market model for Bayesian inference. In *Proceedings of the Twelfth international conference on Uncertainty in artificial intelligence*, pages 405–413. Morgan Kaufmann Publishers Inc.
- Premachandra, M. and Reid, M. D. (2013). Aggregating predictions via sequential mini-trading. In *the 5th Asian Conference on Machine Learning (ACML 2013)*.
- Sethi, R. and Wortman Vaughan, J. (2013). Belief aggregation with automated market makers. Technical report, Microsoft Research New York.
- Shalev-Shwartz, S. and Singer, Y. (2007). Convex repeated games and fenchel duality. In Schölkopf, B., Platt, J., and Hoffman, T., editors, *Advances in Neural Information Processing Systems 19*, pages 1265–1272, Cambridge, MA. MIT Press.
- Storkey, A. (2011). Machine learning markets. In *JMLR Workshop and Conference Proceedings*, volume 15, pages 716–724.
- Storkey, A., Millin, J., and Geras, K. (2012). Isoelastic agents and wealth updates in machine learning markets. In *the 29th International Conference on Machine Learning (ICML 2012)*.
- Von Neumann, J. and Morgenstern, O. (2007). *Theory of games and economic behavior (commemorative edition)*. Princeton university press.
- Wolfers, J. and Zitzewitz, E. (2004). Prediction markets. *The Journal of Economic Perspectives*, 18(2):107–126.
- Yaari, M. E. (1987). The dual theory of choice under risk. *Econometrica*, 55(1):95–115.